arXiv:1411.1690v2 [stat.ML] 9 Mar 2015

Sublinear-Time Approximate MCMC Transitions for Probabilistic Programs

Yutian Chen[1], Vikash Mansinghka[2] and Zoubin Ghahramani[1]

[1]Machine Learning Group, Department of Engineering, University of Cambridge
[2]Computer Science and Artificial Intelligence Laboratory, MIT

Abstract

Probabilistic programming languages can simplify the development of machine learning techniques, but only if inference is sufficiently scalable. Unfortunately, Bayesian parameter estimation for highly coupled models such as regressions and state-space models still scales poorly; each MCMC transition takes linear time in the number of observations. This paper describes a sublinear-time algorithm for making Metropolis-Hastings (MH) updates to latent variables in probabilistic programs. The approach generalizes recently introduced approximate MH techniques: instead of subsampling data items assumed to be independent, it subsamples edges in a dynamically constructed graphical model. It thus applies to a broader class of problems and interoperates with other general-purpose inference techniques. Empirical results, including confirmation of sublinear per-transition scaling, are presented for Bayesian logistic regression, nonlinear classification via joint Dirichlet process mixtures, and parameter estimation for stochastic volatility models (with state estimation via particle MCMC). All three applications use the same implementation, and each requires under 20 lines of probabilistic code.

1 Introduction

Machine learning methods can be difficult and time-consuming to design and implement. Probabilistic programming languages have the potential to mitigate these difficulties and make previously impractical approaches more manageable by formalizing and automating key aspects of modeling and inference. For example, the Stan platform can compactly represent hierarchical Bayesian models with complex structure and provide automatic HMC inference that has been successfully applied to problems with thousands of parameters (Stan Development Team, 2014). Similarly, the Picture language makes it possible to specify and solve challenging 3D vision problems (Kulkarni et al.). Other languages such as Church (Goodman et al., 2008), BLOG (Russell and Li, 2013), Venture (Mansinghka et al., 2014) and Figaro (Pfeffer, 2009) are targeted at a broader class of modeling and inference problems. All these languages go beyond standard toolkits for probabilistic graphical models in important ways.

Probabilistic programming thus provides an appealing setting in which to develop inference algorithms, but scalability is a significant challenge. Many probabilistic programming systems exhibit quadratic scaling on problems such as topic modeling Blei et al. (2003) and nonparametric mixture modeling (Neal, 2000); for example, the Bher transformational compiler (Goodman et al., 2008) requires the entire model to be resimulated for each single-site MH transition. The asymptotic scaling of individual inference steps has been the focus of several research efforts. The Venture and BLOG languages overcome this by tracking de-

pendencies between random choices, using the resulting factorization of the joint probability density over all variables to recover linear scaling in many common cases. The probabilistic execution trace (PET) graph that tracks dependencies in Venture also handles exchangeable coupling and thus supports algorithms based on $O(1)$ updates to sufficient statistics. Other approaches based on program analysis are also being explored (Yang et al., 2014). The scaling problems are particularly severe in highly coupled parameter estimation problems, such as those arising in regression and state-space modeling. For these problems, there are no symmetries to exploit, so each MH update for a parameter requires $O(N)$ time.

This paper describes a sublinear-time algorithm for performing approximate MH transitions to latent variables in probabilistic programs with $O(N)$ outgoing dependencies. This algorithm generalizes ideas from recent work on approximate transition operators (Singh et al., 2012; Korattikara et al., 2014; Bardenet et al., 2014): instead of subsampling data items, it subsamples edges in a dynamically constructed graphical model, stochastically ignoring dependencies. The proposed algorithm can be interleaved with state-of-the-art general-purpose inference algorithms for probabilistic programs and thus applies to problems with widely varying structures.

This paper contains three contributions. The first is the algorithm, including its integration into an inference programming language. The second is a new proof of ergodicity for the approximate Markov chain under milder conditions than Korattikara et al. (2014), showing that the bias vanishes as the controlling parameter approaches 0. The third is an empirical demonstration of efficacy and broad applicability, via applications to parametric regression, nonparametric Bayesian mixtures of experts, and state-space models.

2 Background on Inference in Probabilistic Programs

Here we define the key terms relevant to the subsampled inference technique proposed in this paper. From the perspective of this paper, the function of the background material is to exploit factorization in the joint density of program executions so that individual MH transitions do not have to traverse parts of the execution that have not changed. Due to space constraints, we cannot give a detailed or formal description of complete inference procedures for probabilistic programs. Readers are referred to Mansinghka et al. (2014) for more details.

Probabilistic programming provides a convenient way for users to specify any probabilistic model with a potentially unbounded number of random variables in a formal language and conduct inference conditioned on observed data without implementing inference algorithms by themselves. A probabilistic program describes the generative process of the model with both deterministic or stochastic computations. A typical stochastic computation is to draw a random sample from some distribution. One execution of a program generates one realization of a probabilistic model. Specifically, the program returns a sequence of outputs in the order of execution, one from each computation, denoted by $x_1, x_2, \ldots$. We use $\mathbf{x}_{i:j}$ as a shorthand for $x_i, \ldots, x_j$. The value of x_i depends on its history $\mathbf{x}_{1:i-1}$. So does the existence of one variable itself in the model if the computation appears in a control flow statement.

An example of a probabilistic program in a Lisp style language syntax is shown in the top of Fig. 1. `assume` statements specify the generative model and `observe` statements specify observations, i.e. constraints on random variables. The program first draws a Bernoulli random variable b, assigns the value of variable μ with 1 if $b = True$ or a Gamma random variable otherwise. It then draws a normal random variable y with mean equal to μ and specifies an observation 10. `bernoulli`, `gamma` and `normal` are stochastic computations, and `if` is a deterministic control flow computation. The execution of a program can be represented as a directed graph defined as follows:

Definition 1. A *probabilistic execution trace* (PET or trace) is a directed graph, $\rho := (V, E_\mathrm{e} \cup E_\mathrm{s})$, representing a single run of a probabilistic program, where the set of nodes $V := \{i\}$ represents all executed computations, $E_\mathrm{e} := \{(i,j) : x_i$ is the last variable in $\mathbf{x}_{1:j-1}$ whose value determines the existence of $x_j\}$ represents all *existential* dependencies,

```
[assume b (bernoulli 0.5)]
[assume mu (if b 1 (gamma 1 1))]
[assume y (normal mu 0.1)]
[observe y 10.0]
```

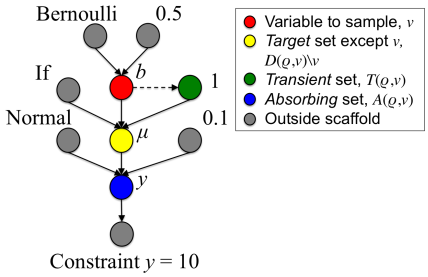

Figure 1: Example of a probabilistic program and its PET when $b = True$. Solid and dashed arrows represent edges in E_s and E_e respectively. Colored nodes specify the scaffold for sampling b which is defined later in Sec. 2.1.

and $E_\text{s} := \{(x_i, x_j) : i < j, x_i \not\perp x_j | \mathbf{x}_{1:j-1\setminus i}\}$ represents all *statistical* dependencies between the value of two nodes.

E_s includes both deterministic and stochastic dependencies while E_e determines the existence of nodes in V. Fig. 1 shows a trace where $b = True$. Because the `if` statement chooses the first branch, the computation of sampling a Gamma random variable, `(gamma 1 1)`, is not executed and therefore does not exist in the trace.

Given V, which depends on E_e, and their statistical dependencies E_s, the probability of generating the trace ρ has the same factorization as a Bayesian network of (V, E_s):

$$p(\rho) = \prod_{n \in V} p(x_n | \text{Par}_\rho(n)) \tag{1}$$

where $\text{Par}_\rho(n) := \{x_{n'} : n' \in V, (n', n) \in E_\text{s}(\rho)\}$ is the parent set of node n in trace ρ.

2.1 Metropolis-Hastings Sampling Algorithm on PETs

Now we describe how to do inference for a probabilistic program with a single-site MH sampling algorithm and its computational complexity. Denote the variable to sample by $x_v, v \in V$. After proposing a new value for x_v, to maintain consistency of the trace we have to propagate the change throughout the descendants of v whose value or existence depends on x_v deterministically.

Definition 2. The *target* set, $D(\rho, v)$, is the set of v and all of its descendants whose computation is always executed for any value of x_v but their value depends on x_v deterministically.

Definition 3. The *transient* set, $T(\rho, v)$, is the set of all the variables whose existence depends on the value of $D(\rho, v)$, $T(\rho, v) := \{n \in V : \exists n' \in D(v, \rho), s.t.\ (n', n) \in E_e\}$.

The change of x_v can be absorbed by children of D or T that represent random computations in the sense that we do not need to propose a new value for those nodes (e.g. node y in Fig. 1) to maintain the consistency of a trace. However, to evaluate the probability of the new trace, we still need to recompute the probability of generating those nodes with the new value of x_v.

Definition 4. The *absorbing* set is $A(\rho, v) := \{n \in V \setminus (D \cup T) : \exists n' \in D \cup T, s.t.\ n' \in \text{Par}_\rho(n)\}$.

For a probabilistic program that describes a regular Bayesian network model without existential and deterministic dependencies, we have the following relationships:

$$D(\rho, v) = \{v\}, \quad T(\rho, v) = \emptyset, \quad A(\rho, v) = \text{Children}(v)$$

Lastly, a scaffold is defined as the union of all the three sets:

Definition 5. The *scaffold* of v in trace ρ is the set of nodes: $s(\rho, v) := D(\rho, v) \cup T(\rho, v) \cup A(\rho, v)$.

One may observe that the scaffold is the set of all the variables whose conditional distribution in Eq. 1 may change with a new proposal for x_v. Particularly when the value of v does not affect the structure of ρ, i.e. $T(\rho, v) = \emptyset$, the union of $s(\rho, v)$ and its parent set excluding D, $(\cup_{n \in s}\{\text{Par}_\rho(n)\} \cup s) \backslash D$, is equivalent to the Markov blanket of v in the corresponding Bayesian network (V, E_s).

Now we have defined all the ingredients to describe the MH algorithm on a PET. Alg. 1 describes the steps to sample variable x_v given a proposal distribution q. A new trace is proposed in step 2-4 and accepted/rejected in the remaining steps. The proposal distribution can be the prior distribution from the program or a customized distribution provided by a user.

Algorithm 1 Metropolis-Hastings Algorithm with Scaffolds

1: **procedure** MH(v, ρ, q)
2: Construct $s(\rho, v)$.
3: *detach*: removes $D \cup T$ from ρ.
4: *regenerate*: proposes new values for D with q and generates a new transient set, T', depending on the new values.
5: Compute acceptance probability P_a.
6: Sample $u \sim \text{Uniform}[0, 1]$.
7: **if** $u \leq P_a$ **then**
8: Accept the new trace.
9: **else**
10: *detach* $D \cup T'$ and restore old values of $D \cup T$.
11: **end if**
12: **end procedure**

Denote the proposed trace with ρ' and the new scaffold with s'. As the structure of ρ does not change except for the transient set, we get $s' = D \cup T' \cup A$. The reverse proposal of MH includes *detaching* $D \cup T'$ from ρ', and *regenerating* old values of D and T. The acceptance probability in Step 5 is then computed as

$$P_a(\rho, \rho') = \min\left\{1, \frac{p(\rho')\prod_{n \in D \cup T} q(x_n|\rho)}{p(\rho)\prod_{n \in D \cup T'} q(x'_n|\rho')}\right\} \tag{2}$$

$$= \min\left\{1, \prod_{n \in D} \frac{p(x'_n|\text{Par}_{\rho'}(n))}{q(x'_n|\rho')} \frac{q(x_n|\rho)}{p(x_n|\text{Par}_\rho(n))} \right.$$
$$\prod_{n \in A} \frac{p(x_n|\text{Par}_{\rho'}(n))}{p(x_n|\text{Par}_\rho(n))} \prod_{n \in T} \frac{q(x_n|\rho)}{p(x_n|\text{Par}_\rho(n))}$$
$$\left. \prod_{n \in T'} \frac{p(x'_n|\text{Par}_{\rho'}(n))}{q(x'_n|\rho')} \right\} \tag{3}$$

where we plugged in Eq. 1 in the second line. $q(x'_n|\rho')$ denotes the proposal distribution of node n given all the nodes in the new trace ρ' that have been generated before n. Any factors for nodes outside $s \cup s' (= D \cup A \cup T \cup T')$ are canceled, consistent with our previous observation about the scaffold. The factors in Eq. 3 depending on ρ' are computed on the fly during Step 4 and the factors depending on ρ are computed during Step 3. For any node $n \in s \cup s'$, we denote by w_n the product of all the factors in Eq. 3 indexed by n. Then we can simplify the acceptance probability as

$$P_a(\rho, \rho') = \min\left\{1, \prod_{n \in s \cup s'} w_n\right\} \tag{4}$$

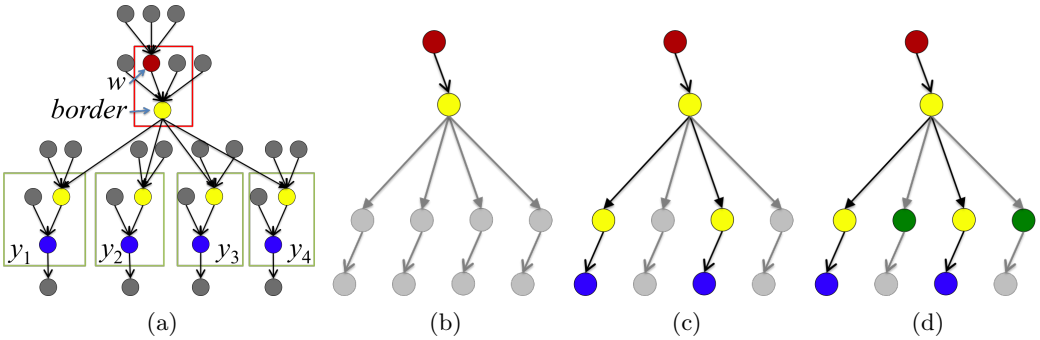

Figure 2: (a) The PET of a Bayesian logistic regression model with 4 observations, colored to encode the scaffold for the weights $\mathbf{w}$; the probabilistic program is shown in Fig. 3. The global section is notated with a red box and local sections with green boxes. (b, c) Subsampled scaffold with no and two local sections. (d) Nodes whose values are stale after a subsampled transition. See Sec. 3.5 for details.

2.2 Computational Complexity of MH Transitions

Eq. 4 makes it clear that the running time of one MH transition scales linearly with the size of the scaffold. For a *local* random variable with a constant size of the scaffold/Markov blanket, it takes $O(1)$ time to run the MH transition. However, for a *global* random variable, e.g. a parameter coupled with all the observations, the size of the scaffold scales with the size of the entire data set. Therefore, in a model with D global variables and N observations, it takes $O(DN)$ time to run MH for just one sweep of all the global variables. That could be unacceptably slow for a problem with a large data set. Given a limited budget on the computational time, a slow transition operator leads to a small number of samples and therefore a large variance in the Monte Carlo estimate to any quantities of interest.

3 Sublinear Time Inference via Subsampling Scaffolds

This paper presents an approximate MH transition for a global variable whose runtime scales sublinearly with the size of data, $o(N)$. We show that the approximate Markov chain is still uniformly ergodic under mild conditions and the induced bias of the stationary distributed is controllable. As we can now generate more samples given the same computational budget, it is possible to improve the accuracy of the MCMC estimate by trading off some bias for a much smaller variance, as confirmed in the experimental section.

3.1 A Global/Local Partition of the Scaffold

We start with an observation of the scaffold for global variables. Fig. 2a shows an example of the scaffold for sampling the weights of a Bayesian logistic regression model with the program specified in the first 7 lines of Fig. 3. The scaffold $s(\rho, v)$ can be partitioned into two parts:

1. A global section, denoted as *global*, that contains v and dependencies that do not scale with N.
2. A set of N local sections, $\{local_n\}_{n=1}^{N}$, which share a similar structure and represent the N dependencies.

The size of the scaffold grows with the number of local sections while the size of each section remains constant.

Our algorithm applies to scaffolds of this structure, where in addition we assume all the local dependencies are connected from v through a common node with a single link, which is v itself or other node in $D(\rho, v)$. Under this assumptions, the common node, e.g. the top yellow node in Fig. 2b, is the border that separates *global* from *local* sections and has N

children. We also assume that $T(\rho, v) = \emptyset$, i.e. making proposals for x_v does not change the structure of the PET. Note that this does *not* exclude models with dynamically evolving structure; only approximate transitions are prohibited from introducing structural changes. These are mild restrictions, still admitting a broad class of probabilistic programs.

Formally, we provide constructive definitions for the terms above

Definition 6. The *border* node, $b(s, v)$, of a scaffold $s(\rho, v)$, is the first descendant of v with multiple branches in s.

Definition 7. The *global* section of a scaffold $s(\rho, v)$ is: $global(\rho, v) := s(\rho, v) \backslash desendants(\rho, b)$.

Definition 8 The *local* section of a scaffold $s(\rho, v)$ associated with i'th child of b, c_i, is: $local_i := s \cap (c_i \cup desendants(\rho, c_i))$.

The scaffold is then partitioned into one *global* and N *local* sections mutually exclusively.

3.2 Austerity MCMC for Probabilistic Programs with Partitioned Scaffolds

Korattikara et al. (2014) proposed the austerity idea, a sublinear-time algorithm to incrementally approximate the acceptance probability until the approximation error is under a tolerance level. They applied the algorithm to the posterior inference problem for Bayesian models with N *iid* observations. We adopt this idea for the inference problem on PETs by proposing a sublinear-time algorithm for approximate MH transitions on partitioned scaffolds.

Given a partition of the scaffold, we can factorize the acceptance probability in Eq. 4 as

$$P_a = \min \left[1, \left(\prod_{n \in global} w_n \right) \prod_{i=1}^{N} \left(\prod_{n \in local_i} w_n \right) \right] \tag{5}$$

After drawing a uniform random variable u, the inequality $u < P_a$ is equivalent to $\mu_0 < \mu$ with μ_0 and μ defined as follows:

$$\mu_0 = \frac{1}{N}(\log u - \sum_{n \in global} \log w_n), \mu = \frac{1}{N} \sum_{i=1}^{N} l_i, \tag{6}$$

where $l_i = \sum_{n \in local_i} \log w_n$. The term μ_0 is usually cheap to compute, while μ is an expensive average over N terms when N is large. We follow Korattikara et al. (2014) and reformulate the decision problem as a statistical test with the hypothesis $H_1 : \mu > \mu_0$, and $H_2 : \mu < \mu_0$. Given μ_0, the set $\{l_i\}$ and a tolerance level ε, we conduct the hypothesis testing by iteratively sampling mini-batches of m local sections, and updating the estimate of μ until a high confident result is reached. The algorithm is given in Alg. 2. The last term in $\sqrt{\cdot}$ in step 7 is called the finite population correction factor introduced by sampling without replacement. In step 8, if $s_l = 0$, we do not run the t-test but continue to draw another subset until $n = N$. That prevents a false decision in early stages when a small subset of $\{l_i\}$ includes all equal numbers by chance. For problems with a large data set, the approximate MH method requires only a small fraction of the data to make a high quality decision, and therefore can significantly speed up the MH algorithm.

By comparing the definition of μ_0 and μ with those in Korattikara et al. (2014), one may find out that μ_0 in both methods contains the same prior and proposal distribution for x_v. However, the quantity μ in our method has a more general meaning. In Korattikara et al. (2014), each l_i refers to the log-ratio of the likelihood of one *iid* data point, while in our PET setting each l_i is the product of w's associated with a local partition. They do not have to be observed data, and can even have strong dependence between each other; our third experiment illustrates this case.

Remark. The same error analysis in Korattikara et al. (2014) applies to our algorithm for PETs even at the presence of dependencies between local sections in the model distribution.

Algorithm 2 Sequential Test for MH Decison

procedure SEQUENTIALTEST(μ_0, set $\mathcal{X} = \{l_i\}$ of a size N, mini-batch size m, error tolerance ε)
  Set of observed data $\mathcal{Y} \leftarrow \emptyset$, $n \leftarrow 0$.
  repeat
    Draw a subset of m l_i's, $X_m \subseteq \mathcal{X}$
    $\mathcal{X} \leftarrow \mathcal{X} \backslash X_m$, $\mathcal{Y} \leftarrow \mathcal{Y} \cup X_m$, $n \leftarrow n + m$
    Estimate mean, $\hat{\mu}$, and std, s_l, of $\mathcal{Y}$.
    std of $\hat{\mu}$ is $s \leftarrow \frac{s_l}{\sqrt{n}} \sqrt{1 - (n-1)/(N-1)}$
  until p-value of $t_{n-1}(|\frac{\hat{\mu} - \mu_0}{s}|)$ is lower than ε
  Accept H_1 **if** $\hat{\mu} > \mu_0$ **else** H_2
end procedure

That remark follows the fact that the sequential test procedure in Alg. 2 exploits the statistical property of the randomly subsampled l_i when the entire set $\{l_i\}$ is given. *Conditioned* on $\{l_i\}$, the samples randomly drawn without replacement are independent with each other except for the correlation introduced by the "without replacement" part. Therefore, error induced by the sequential test does not depend on whether there exist dependencies among l_i's in the prior distribution defined by the probabilistic program.

3.3 Robustness for General Problems

It is worth noting that the error of the sequential test depends on the *marginal* distribution of $\{l_i\}$'s. This is because our algorithm is based on student-t test. When there exist many outliers with extremely large values, as demonstrated by a synthetic counter-example in Bardenet et al. (2014), the central limit theorem breaks down on the subset of $\{l_i\}$. In that case, our approximate MH transition might lead to noticeable bias in the stationary distribution. Despite the potential failure in the synthetic example, the austerity MH algorithm has shown robust performance in all the real-data experiments from both Korattikara et al. (2014) and this paper with a proper size of a mini-batch.

Our software can provide a normality test for the distribution of the estimated mean $\hat{\mu}$ in trial runs and produce an auto-generated comparison between the performance of the approximate MH and regular inference. One advantage of a probabilistic programming implementation is that users do not need to implement safeguards themselves. In practice, we suggest applying our proposed algorithm on variables with a single type of dependence to reduce the risk of outliers.

Additionally, we prove that even when the condition for the central limit theorem does not hold, the bias of the approximate Markov chain still diminishes as the tolerance parameter ε in Alg. 2 approaches 0 under mild conditions.

Let θ and θ^* be the current and proposed value of x_v, m be the size of a mini-batch, and $P_{a,\varepsilon}(\theta, \theta^*)$ be the acceptance probability using the sequential test with ε. We have the following theorem with proof in the supplementary.

Theorem 1. *If either (1) the domain of θ, Θ, is compact and the likelihood function is continuous w.r.t. θ, or (2) Θ, is a finite set, then there exists a function $\delta(\varepsilon)$ such that $|P_{a,\varepsilon}(\theta, \theta^*) - P_a(\theta, \theta^*)| \leq \delta(\varepsilon), \forall \theta, \theta^* \in \Theta$, and $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$.*

This theorem shows a uniform convergence of the approximate MH transition to the exact transition. The compactness condition might be violated if θ has an unbounded domain. However, its impact to the bias of the stationary distribution could be negligible when the actual samples of θ reside in a finite region for any sufficiently small ε. Alternatively, one can use a truncated prior distribution for θ.

We can further prove the uniform ergodicity of our approximate Markov chain as follows.

Corollary 2. *If either condition in Thm. 1 holds and the exact Markov chain satisfies the regularity assumptions 1-3 in Sec. 3.2 of Pillai and Smith (2014), the Markov chain of the*

approximate MH algorithm is uniformly ergodic for any sufficiently small ε. Its stationary distribution approaches the target distribution as $\varepsilon \to 0$.

3.4 Implementation with Scaffold Subsampling

In order to obtain sublinear running time, we should avoid any operations with a $O(N)$ complexity, including the construction of the entire scaffold. Our proposed algorithm interleaves the scaffold construction with the *detach* & *regenerate* operations. This way the local sections of the scaffold will not be created until the sequential test requires more data to improve its estimate. Let Detach&Regen(A) be the procedure to detach and regenerate a set of nodes $A \subseteq s$ and return $\sum_{n \in A} \log w_n$. Our subsampled MH algorithm is given in Alg. 3. Fig. 2(b, c) illustrate the state after step 4 and after two local partitions are constructed respectively.

Algorithm 3 Sublinear-Time Metropolis-Hastings Algorithm with Scaffolds

1: **procedure** SUBSAMPLEDMH(v, ρ, q, m, ε)
2: Sample $u \sim$ Uniform$[0, 1]$.
3: Find $b(s, v)$ as Def. 6, $N \leftarrow \#(b$'s children)
4: Construct *global* as Def. 7 by contructing a regular $s(\rho, v)$ without reaching beyond b
5: $\sum_{n \in global} \log w_n \leftarrow$ Detach&Regen(*global*)
6: Compute μ_0 with Eq. 6, $\mathcal{Y} \leftarrow \emptyset$, $n \leftarrow 0$
7: **repeat**
8: Sample m children of b's w/o replacement, $\{c_i\}$.
9: **for all** c_i **do**
10: Construct $local_i$ as Def. 8 by resuming the construction of $s(\rho, v)$ from c_i
11: $l_i \leftarrow$ Detach&Regen($local_i$)
12: **end for**
13: $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{l_i\}$, update n, $\hat{\mu}$, and s as in Alg. 2
14: **until** The p-value falls below ε as in Alg. 2
15: **if** $\hat{\mu} > \mu_0$ **then**
16: Accept the new trace.
17: **else**
18: *detach* and restore old values of *global*.
19: **end if**
20: **end procedure**

3.5 Updating stale nodes on demand

When a proposed move is accepted, all the deterministic nodes, $D(\rho, v)$, should be updated accordingly in the standard MH algorithm. However, with the subsampling approach it is not guaranteed to update all *local* sections if some of them have not been constructed. This may leave some nodes with old values and break the deterministic dependencies. The green nodes in Fig. 2d are examples of the stale nodes. In order to solve the broken consistency while retaining sublinear runtime, we proposed a lazy updating approach: instead of updating deterministic nodes in local sections immediately after a proposal is accepted/reject, we leave the stale deterministic nodes as is but whenever it is to be accessed later, we first update their values and then process the node.

4 Experiments

We implemented our algorithm in a lightweight research variant of the Venture probabilistic programming system, written in unoptimized Python, and applied it to multiple problems. Asymptotic scaling and relative comparisons between standard and subsampled MH are meaningful, but absolute runtimes reflect large constant factor overheads that are straightforward to reduce.

Table 1: Overview of models used in experiments, and scaling parameters for exact MH. $N_k := \sum_{i=1}^{N} \mathbb{I}[z_i = k]$.

Model	Domain of Sublinear MH	Scaling
BayesLR	$\mathbf{w} \sim p(\mathbf{w}) \prod_{i=1}^{N} \text{Logit}(y\|\mathbf{x}_i, \mathbf{w})$	N
JointDPM	$\mathbf{w}_k \sim p(\mathbf{w}_k) \prod_{i:z_i=k} \text{Logit}(y\|\mathbf{x}_i, \mathbf{w})$	N_k
SV	$\phi/\sigma \sim p(\phi/\sigma) \prod_{t=1}^{T} \mathcal{N}(h_t\|\phi h_{t-1}, \sigma^2)$	T

```
[assume w (scope_include 'w 0
           (multivariate_normal {mu} {Sig}))]
[assume y_x (lambda (x) (bernoulli (linear_logistic w x)))]

for n in 1...N: #load data
  # features X[n] = [x0_n x1_n ...], class Y[n]
  [observe (y_x {X[n]}) {Y[n]}]

# do T iterations of subsampled MH with a Gaussian drift
# proposal of bandwidth sig
[infer (subsampled_mh w all {nbatch} {eps} 'drift {sig} {T}]
[infer (mh w all {T})] #standard MH provided for reference
```

Figure 3: Probabilistic program for the Bayesian logistic regression model, data, and subsampled inference scheme. {} denotes external parameters. (`scope_include` 'w 0 x) is a deterministic computation that adds label w to x for the `infer` statement to specify which variable to sample.

4.1 Bayesian Logistic Regression

We first demonstrate that our general-purpose implementation recovers the time/accuracy advantages of the custom implementation of austery for Bayesian logistic regression from Korattikara et al. (2014). Our Bayesian logistic regression model uses a standard isotropic Gaussian prior on the weights:

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, 0.1\mathbf{I}_D), \quad y_i \overset{\text{iid}}{\sim} \text{Logit}(y|\mathbf{x}_i, \mathbf{w}) \tag{7}$$

We evaluate performance on a classification task on the MNIST digit image data. We train on 12214 images of '7' and '9', each transformed to a 50-dimensional feature vector via normalization and principal component analysis. We evaluate predictive accuracy versus computation time on 2037 test images. The Venture programs for the model and inference is given in Fig. 3. We use the same parameter and data settings as in Korattikara et al. (2014), and run the inference algorithm with the same random walk proposal distribution. We use a smaller mini-batch size of 100. The risk of predictive mean over an extensive run over 50 hours are given in Fig. 4. See Korattikara et al. (2014) for the definition of risk. The predictive performance of subsampled MH increases significantly faster than the standard MH. It can make more than one order of magnitude more transitions, and takes 5 hours to reach the risk achieved by standard MH after 50 hours.

4.1.1 Sublinearity

Here we empirically study the asymptotic runtime of one MH iteration with the proposed algorithm for this model. We run the same algorithm on a synthetic dataset with two input features and one binary label so that we can easily control the amount of labeled data. The dataset is illustrated in Fig. 5a. We fix $\epsilon = 0.01$, a mini-batch of 100, and a standard deviation of the proposal distribution of 0.1, and vary the number of data points. The blue line in Fig. 5b shows the expected number of subsampled data points at a particular iteration for different sizes of the dataset in the log-log scale. The expectation is computed theoretically with Eqn. 19 in Korattikara et al. (2014). We use the same current and proposed parameter value for all dataset sizes in this plot. The green line shows the empirical average number of subsampled data points over 300 iterations. Fig. 5c shows the average

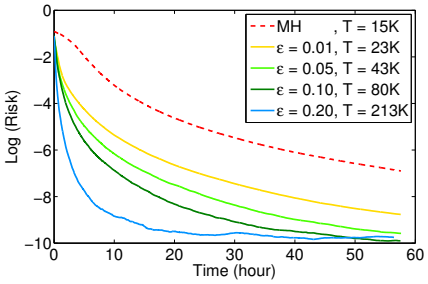

Figure 4: Risk of predictive mean vs. running time for Bayesian logistic regression on MNIST. T in the legend denotes the number of samples drawn in about 60 hours.

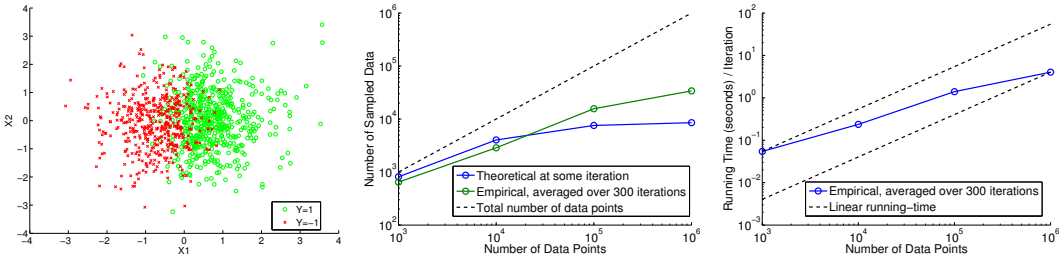

(a) (b) (c)

Figure 5: Sublinear behavior on a synthetic data set. (a) Training data. (b) Number of subsampled data points per iteration in log-log scale. (c) Running time per iteration in log-log scale.

running time per iteration with dashed lines indicating the behavior of a linear algorithm. Both the number of subsampled data points and running time confirm that our proposed subsample MH exhibits sublinear-linear asymptotic scaling for *global* parameters given a fixed proposal distribution.

4.2 Joint Dirichlet Process Mixture Models

We also assess performance on a joint Dirichlet process mixture (JointDPM) model (Wade et al., 2014), a flexible nonlinear classifier that combines logistic regressions using nonparametric Bayes. This model uses a Dirichlet process mixture of Gaussians to model input features, where each component has a distinct set of logistic regression weights to classify the vectors it contains. More formally, we have

$$(\mathbf{x}_i, y_i)|P \overset{\text{iid}}{\sim} f(\mathbf{x}, y|P), \quad P \sim \text{DP}(\alpha P_0) \tag{8}$$

where P_0 is the base measure. Every sample of P is a countable mixture model

$$f(\mathbf{x}, y|P) = \sum_{k=1}^{\infty} \pi_k \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \Sigma_k)\text{Logit}(y|\mathbf{x}, \mathbf{w}_k) \tag{9}$$

with each normal distribution parameter pair $(\boldsymbol{\mu}_k, \Sigma_k)$ sampled from a conjugate normal-inverse-Wishart prior. The regression parameter $\mathbf{w}_k$ is sampled from an isotropic Gaussian prior as in the Bayesian logistic regression model. The graphical model is shown in Fig. 8a. We can write the model in under 20 lines of probabilistic code (Fig. 7), collapsing the component models by marginalizing out (μ_k, Σ_k) and collapsing the DP into a CRP. A trace fragment containing two scaffolds, each containing the regression weights for one cluster, is shown in Fig. 6a. In this model, the number of simultaneous instantiations of the austerity scheme is an object of inference.

We use subsampled MH to accelerate inference over the parameters w_k for each cluster's logistic regression model. The CRP hyperparameter α and N component assignment variables z_i must also be inferred, but as inference for these variables already requires constant

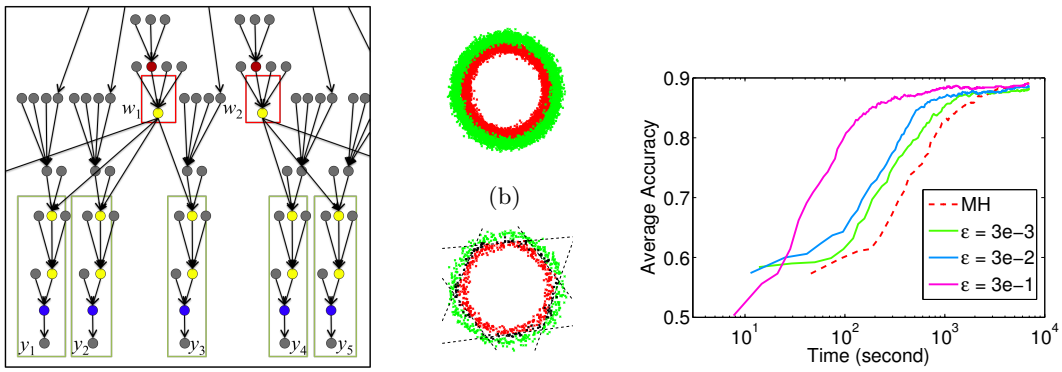

(a) (b) (c) (d)

Figure 6: Joint Dirichlet Process mixture model evaluated on synthetic data. (a) part of PET with scaffolds for two weight parameters. (b) training data. (c) prediction on test data with $\varepsilon = 0.3$ after 2 hours. 6 clusters are found with decision boundaries (dashed lines) and misclassified points (black dots). (d) Predictions accuracy vs running time in log domain.

```
# fixed: dimensionality D, hypers mu_w sig_w m0 k0 v0 S0
[assume alpha (scope_include 'hypers 0 (gamma 1 1))]
[assume crp   (make_crp alpha)]
[assume z     (mem (lambda (i) (scope_include 'z i (crp))))]
[assume w (mem (lambda (z) (scope_include 'w z
        (multivariate_normal mu_w sig_w))))]
[assume c (mem (lambda (z)
        (make_collapsed_multivariate_normal m0 k0 v0 S0)))]
[assume x (lambda (i) ((c (z i))))]
[assume y (lambda (i x)
           (bernoulli (linear_logistic (w (z i)) x)))]
for i in 1...X: # load data
    [observe (x i) {X[i]}] #X[i] is ith feature vector
    [observe (y i) {Y[i]}] #Y[i] is ith class label

# T steps of MH for hyperparams, single-site gibbs for z,
# subsampled MH over weights for a randomly chosen expert
[infer (cycle {T} ((mh alpha all 1)
                   (gibbs z one step_z)
                   (subsampled_mh w one {Nbatch} {eps}
                                  'drift {sigma} 1)) 1)]
```

```
[assume sig (scope_include 'sig 0 (sqrt (inv_gamma 5 0.05)))]
[assume phi (scope_include 'phi (beta 5 1))]
[assume h   (mem (lambda (t) (scope_include 'h t
              (if (<= t 0) 0
                  (normal (* phi (h (- t 1))) sig)))))]
[assume x (lambda (t) (normal 0 (/ (h t) 2)))]
for t in 1...T:
  [observe (x t) {X[t]}] #X[t] is the observation at time t

# state estimation via particle gibbs over subsequences of
# length L
for n in 1...N-L:
    [infer (pgibbs h (ordered_range h h+L) P 1)]
# subsampled MH inference on the parameters
[infer (cycle ((subsampled_mh sig 0 {Nbatch} {eps} 1)
               (subsampled_mh phi 0 {Nbatch} {eps} 1)) 1)]
```

Figure 7: Probabilistic program containing model and inference scheme for (top) joint DPM in Sec. 4.2 and (bottom) stochastic volatility model in Sec. 4.3.

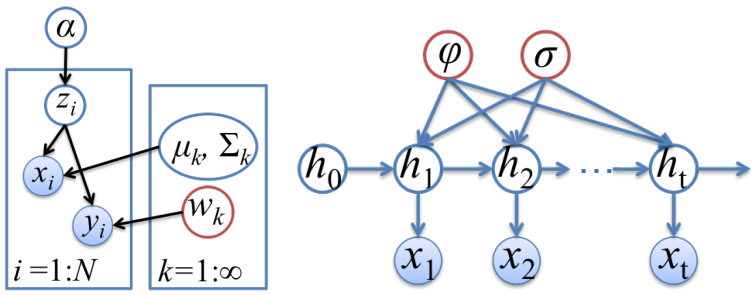

(a) (b)

Figure 8: Graphical model of (a) Joint DPM (b) stochastic volatility model. Subsampled MH applies to red nodes. The algorithm from this paper applies to problems with diverse dependency structures.

time due to properties of the PET, approximate transitions are only used for the w_ks. We set the parameters of our inference program to allocate roughly equivalent computation time to sampling w_k and to a series of transitions to randomly chosen z_i's; the balance is struck using the `step_z` parameter (see Fig. 7). Because subsampled MH runs faster than MH, the w's are updated more frequently than standard MH to maintain the balance.

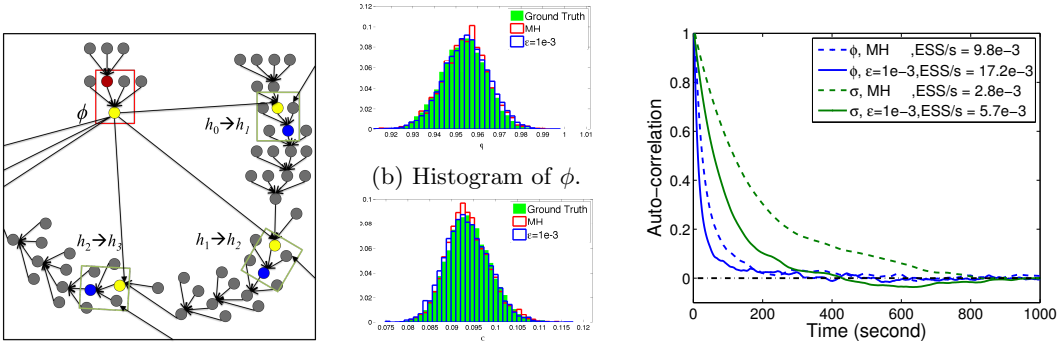

(a) (b) Histogram of ϕ. (c) Histogram of σ. (d)

Figure 9: Stochastic Volatility Model. (a): part of PET and the scaffold for sampling ϕ. (b, c): Histogram of the samples of ϕ and σ from ground truth (green), exact MH (red) and subsampled MH with ε=1e-3 (blue). (d): Auto-correlation of samples for exact (dashed) and subsampled MH (solid). ESS per second shown in the legend.

We apply JointDPM to a synthetic data set with 10,000 data points as shown in Fig. 6b. A snapshot of the prediction results on another 1,000 test points and the accuracy of the prediction averaged over the Markov chain are shown in Fig. 6c and 6d. The sublinear algorithm with $\varepsilon = 0.3$ reaches the same accuracy as exact MH in 10x less time.

4.3 Joint Parameter and State Estimation in a State-Space Model

We also applied our implementation to a stochastic volatility model for a time series with two parameters:

$$x_t = \exp(h_t/2)\varepsilon_t,\ h_t \sim \mathcal{N}(\phi h_{t-1}, \sigma^2),\ \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$$

where we set $h_0 = 0$ and assign a Beta(5,1) and an inverse Gamma(5, 0.05) distribution for parameter ϕ and σ^2 respectively. This model is a state-space model and has unknown hidden states as well as unknown parameters. Fig. 8b shows the graphical model. Fig. 9a illustrates the PET of the model with $t = 1, 2, 3$. Note that there are dependencies between the subsampled local partitions in this problem, transitions from h_{t-1} to h_t. The dynamics are sensitive to the precise values of ϕ and σ^2 and also the hidden states, yielding a challenging joint inference problem.

We generate a synthetic data set of 200 series of length 5 with correlation $\phi = 0.95$ and noise $\sigma = 0.1$. We apply particle Gibbs to sample latent states and (subsampled) MH to sample parameters ϕ and σ. Fig. 9b and c show the histogram of the samples after the burn-in period, and Fig. 9d shows the autocorrelation of the samples measured in running time. Due to the high correlation between latent state variables, the overall mixing rate highly depends on the mixing rate in h_t's. Therefore we assign 10 more times of computation time to sample h_t than other variables. We observe that the gain of subsampled MH is not as significant as previous experiments due to the slow mixing of latent states. Nevertheless, subsampled MH still obtains about twice the efficiency of exact MH without introducing significant bias.

5 Discussion

This paper shows that it is feasible to define sublinear approximate MH transitions for highly-coupled variables in general probabilistic programs and integrate them into an inference programming language. The results also suggest that it is useful: a single unoptimized implementation applies to a broad class of problems, going beyond previous approximate MH schemes, and yields significant (2x-10x) improvements in runtime at no cost in accuracy.

The algorithm presented here uses frequentist statistical inference to accelerate Bayesian inference. It is interesting to consider alternatives to Algorithm 1, where the sequential test is replaced with a model-based inference — potentially itself written as a probabilistic program. It may also be interesting to consider probabilistic programming adaptations of other techniques for accelerating inference stochastically ignoring or suppressing dependencies, such as decayed MCMC filtering (Golightly et al., 2014).

Probabilistic programs can represent model classes, datasets, queries and custom inference strategies from many different application domains. Some domains will require faithful, fully Bayesian inference. In these cases, approximate MH may be useful for accelerating burn-in. For other problems, such as parameter estimation for motion models in robotics or machine learning from internet user behavior, speed may be preferred to accuracy, and results from sublinear approximate MH transitions may be adequate on their own. We hope that by integrating sublinear, approximate transitions into a higher-order probabilistic programming system, we have taken a step towards making it more practical to use probabilistic programming in both these kinds of applications.

References

Rémi Bardenet, Arnaud Doucet, and Chris Holmes. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *Proceedings of The 31st International Conference on Machine Learning*, pages 405–413, 2014.

D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

Andrew Golightly, DanielA. Henderson, and Chris Sherlock. Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing*, pages 1–17, 2014. ISSN 0960-3174.

Noah Goodman, Vikash Mansinghka, Daniel Roy, Keith Bonawitz, and Joshua Tenenbaum. Church: a language for generative models. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 220–229, Corvallis, Oregon, 2008. AUAI Press.

Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 181–189, 2014.

Tejas D. Kulkarni, Pushmeet Kohli, Josh Tenenbaum, and Vikash Mansinghka. Picture: An imperative probabilistic programming language for scene perception. Under review.

Vikash Mansinghka, Daniel Selsam, and Yura Perov. Venture: a higher-order probabilistic programming platform with programmable inference. *arXiv preprint arXiv:1404.0099*, 2014.

R.M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9:283–297, 2000.

A Pfeffer. Figaro: An object-oriented probabilistic programming language. Technical report, Charles River Analytics, 2009.

Natesh S Pillai and Aaron Smith. Ergodicity of approximate MCMC chains with applications to large data sets. *arXiv preprint arXiv:1405.0182*, 2014.

Stuart J. Russell and Lei Li. The blog language reference. Technical report, EECS Department, University of California, Berkeley, 2013.

Robert J Serfling. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics*, pages 39–48, 1974.

Sameer Singh, Michael Wick, and Andrew McCallum. Monte Carlo MCMC: efficient inference by approximate sampling. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1104–1113. Association for Computational Linguistics, 2012.

Stan Development Team. *Stan Modeling Language Users Guide and Reference Manual, Version 2.5.0*, 2014. URL `http://mc-stan.org/`.

Sara Wade, David B. Dunson, Sonia Petrone, and Lorenzo Trippa. Improving prediction from Dirichlet process mixtures via enrichment. *Journal of Machine Learning Research*, 15:1041–1071, 2014.

Lingfeng Yang, Pat Hanrahan, and Noah D Goodman. Generating efficient mcmc kernels from probabilistic programs. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 1068–1076, 2014.

A Proof of Theorems 1

Proof. Given the value of (θ, θ^*), if all numbers in the set $\{l_i\}$ all equal, Alg. 2 would always produce the correct decision and therefore $P_a(\theta, \theta^*) = P_{a,\varepsilon}(\theta, \theta^*)$. We only consider the case when all the numbers are not equal from now on. Since the size of the set $\{l_i\}$ is finite, there exists a positive minimal value for any positive estimated standard deviation obtained from a subset of $\{l_i\}$:

$$s_l^*(\theta, \theta^*) \stackrel{\text{def}}{=} \min_{S \subseteq (\{l_i\}): s_l(S) > 0} s_l(S) > 0.$$

In the accept/reject phase of a MH iteration, given the uniform random variable u, the sequential test algorithm will make a wrong decision if at some step in the sequence of tests $\hat{\mu} - \mu_0 > s t_{n-1}^{-1}(1-\varepsilon)$ while the true mean $\mu \leq \mu_0$ or $\hat{\mu} - \mu_0 < -s t_{n-1}^{-1}(1-\varepsilon)$ while $\mu > \mu_0$, where t_n^{-1} is the inverse CDF of student t distribution with a degree of freedom n. Letting $C(\theta, \theta^*) \stackrel{\text{def}}{=} \max_i\{|l_i|\}$, as we do a test only when $s_l > 0$, we can bound the probability of making an error at step t by applying the one side Hoeffding's inequality with adjustment for sampling without replacement (Serfling, 1974) as

$$\begin{aligned}
&P\left(\hat{\mu} - \mu_0 > s t_{n-1}^{-1}(1-\varepsilon)\right) \\
&\leq P\left(\hat{\mu} - \mu > \frac{s_l^*}{\sqrt{n}}\sqrt{1 - \frac{n-1}{N-1}} t_{n-1}^{-1}(1-\varepsilon)\right) \\
&\leq \exp\left(-\frac{1}{2}\left(\frac{s_l^* t_{n-1}^{-1}(1-\varepsilon)}{C}\right)^2\right) \\
&\stackrel{\text{def}}{=} f_1(\varepsilon, n, \theta, \theta^*) \qquad (10)
\end{aligned}$$

when $\mu \leq \mu_0$. In the fist inequality we plugged in the assumption $\mu \leq \mu_0$, the definition of s in Step 8 of Alg. 2 and the lower bound of s_l. In the second inequality we applied the

concentration bound. It is obvious to observe that the error is also upper bounded by f_1 when $\mu > \mu_0$. Notice that $f_1(\varepsilon, n, \theta, \theta^*) \to 0$ as $\varepsilon \to 0, \forall n, \theta, \theta^*$.

We can furthur bound the total probability of making a wrong decision in the sequential test given u, denoted as $\mathcal{E}(u, \varepsilon)$, using the union bound:

$$\begin{aligned}\mathcal{E}(u, \varepsilon, \theta, \theta^*) &= P\left(\cup_t \{\text{wrong decision at step } t\}\right) \\ &\leq \sum_{t=1}^{\lceil N/n \rceil} f_1(\varepsilon, mt, \theta, \theta^*) \stackrel{\text{def}}{=} f_2(\varepsilon, \theta, \theta^*) \end{aligned} \tag{11}$$

We omit the dependency of f_2 on N and m. Again, we have $f_2(\varepsilon, \theta, \theta^*) \to 0$ as $\varepsilon \to 0, \forall \theta, \theta^*$.

Denote the acceptance probability of our approximate Markov chain at (θ, θ^*) as $P_{a,\varepsilon}$ and we can bound its error as

$$\begin{aligned}&|P_{a,\varepsilon}(\theta, \theta^*) - P_a(\theta, \theta^*)| \\ &= \left| \int_0^1 P(A|u) \mathrm{d}u - \int_0^{P_a} \mathrm{d}u \right| \\ &= \left| \int_{P_a}^1 P(A|u) \mathrm{d}u - \int_0^{P_a} (1 - P(A|u)) \mathrm{d}u \right| \\ &\leq \int_0^1 \mathcal{E}(u, \varepsilon, \theta, \theta^*) \mathrm{d}u \leq f_2(\varepsilon, \theta, \theta^*) \end{aligned} \tag{12}$$

where A denotes the event that the sequential test procedure accepts the proposal. So $P(A|u)$ with $u > P_a$ is the probability of accepting a proposal while we should reject it and $1 - P(A|u)$ with $u < P_a$ is the probability of rejecting a proposal when we should accept it.

Now let us consider the two conditions in Theorem 1 separately.

Under condition (1), the value of l_i as a function of (θ, θ^*) is continuous. Therefore, the functions s^* and C are also continuous w.r.t. (θ, θ^*). So is the upper bound $f_2(\theta, \theta^*)$.

Combined with the condition that Θ is compact, we can conclude the proof by claiming that function f_2 will achieve its maximum in the domain of $\Theta \times \Theta$, denoted as $\delta(\varepsilon)$, and because the function f_2 approaches 0 everywhere in the compact set $\Theta \times \Theta$ as $\varepsilon \to 0$, $\delta(\varepsilon)$ also approaches 0 with ε.

Under condition (2), since the domain of θ is finite, there exists a maximum value of $f_2(\varepsilon, \theta, \theta^*)$ over (θ, θ^*) for any ε. Let

$$\delta(\varepsilon) \stackrel{\text{def}}{=} \max_{\theta, \theta^* \in \Theta \times \Theta} f_2(\varepsilon, \theta, \theta^*) \tag{13}$$

Since $\forall \theta, \theta^* \in \Theta \times \Theta$, $f_2(\varepsilon, \theta, \theta^*) \to 0$ as $\varepsilon \to 0$, it follows that $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$. □

B Proof of Corollary 2

Proof. Denote the proposal distribution with $q(\theta, \theta')$. For a variable defined in a compact space Θ with measure Ω, the error of expected rejection probability of the approximate Markov chain is bounded by

$$|P_{r,\varepsilon} - P_r| = \left| \int_{\theta'} (-P_{a,\varepsilon} + P_a) q(\theta, \theta') \mathrm{d}\Omega(\theta') \right| \leq \delta(\varepsilon) \tag{14}$$

The transition kernel of M-H is $\mathcal{T}(\theta, \theta') = P_a q(\theta, \theta') + P_r \delta(\theta, \theta')$, where $\delta(\theta, \theta')$ denotes the Dirac delta function. We can bound the total variation distance between the approximate

M-H kernel and the exact kernel as

$$
\begin{aligned}
&\|\mathcal{T}_\varepsilon(\theta,\cdot)-\mathcal{T}(\theta,\cdot)\|_{\mathrm{TV}} \\
&= \frac{1}{2}\int_{\theta'}\Big|\left(P_{a,\varepsilon}(\theta,\theta^*)-P_a(\theta,\theta')\right)q(\theta,\theta') \\
&\qquad\qquad +\left(P_{r,\varepsilon}-P_r\right)\delta(\theta,\theta')\Big|\mathrm{d}\Omega(\theta') \\
&\le \delta(\varepsilon), \forall\theta\in\Theta
\end{aligned}
\tag{15}
$$

Since $\delta(\varepsilon)\to 0$ as $\varepsilon\to 0$, for any sufficiently small ε we can apply the Lemma 3.6 of Pillai and Smith (2014) to prove the uniform ergodicity and obtain the convergence rate. □